OLCF Seminar Series

MAGMA - a New Generation of Linear Algebra Libraries for GPU and Multicore Architectures

June 16, 2010

2:00 p.m.

Bldg.5100
Room 128

HOST: Oscar Hernandez (oscar@ORNL.GOV

Note: If you would like to request a meeting with the visitor, please contact Sherry Hempfling (hgv@ornl.gov)

Dr. Stanimire Tomov
The University of Tennessee
Center for Information Technology Research
319 Philander P. Claxton Education Building
1122 Volunteer Boulevard
Knoxville, TN 37996-3400

Abstract. This talk will highlight the latest accomplishments in the Matrix Algebra on GPU and Multicore Architectures (MAGMA) project. We use a hybridization methodology that is built on representing linear algebra algorithms as collections of tasks and data dependencies, as well as properly scheduling the tasks' execution over the available multicore and GPU hardware components.

This methodology is applied in MAGMA to develop high-performance fundamental linear algebra routines, such as the one-sided dense matrix factorizations (LU, QR, and Cholesky) and linear solvers, two-sided dense matrix factorizations (bidiagonal, tridiagonal, and Hessenberg reductions) for singular and eigenvalue problems, in addition to iterative linear and eigenvalue solvers. MAGMA is designed to be similar to LAPACK in functionality, data storage, and interface, in order to allow scientists to effortlessly port any of their LAPACK-relying software components to take advantage of the new architectures. Finally, we will present MAGMA developments targeting the Fermi GPU.

BIO: Stanimire (Stan) Tomov is Research Scientist at ICL and Adjunct Assistant Professor in EECS at UTK. He received Ph.D. in Mathematics from TAMU in 2002 and held positions at LLNL and BNL. Stan's research interests are in parallel algorithms, numerical analysis, and high- performance scientific computing. He co-leads the UTK CUDA Center of Excellence on the development of Linear Algebra Libraries for CUDA-based Hybrid Architectures, and in particuler MAGMA, a new generation of linear algebra libraries, extending the sequential LAPACK-style algorithms, for highly parallel, GPU and multicore heterogeneous architectures.